

2 A high-performance network intrusion detection system R. Sekar, Y. Guang, S. Verma, T. Shanbhaq

November 1999 Proceedings of the 6th ACM conference on Computer and communications security CCS '99

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.04 MB) terms

In this paper we present a new approach for network intrusion detection based on concise specifications that characterize normal and abnormal network packet sequences. Our specification language is geared for a robust network intrusion detection by enforcing a

strict type discipline via a combination of static and dynamic type checking. Unlike most previous approaches in network intrusion detection, our approach can easily support new network protocols as information relating to the protoco ...

Adaptive Caches: Effective Shaping of Cache Behavior to Workloads Ranjith Subramanian, Yannis Smaragdakis, Gabriel H. Loh December 2006 Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture MICRO 39

Publisher: IEEE Computer Society

Full text available: 🔁 pdf(1.41 MB) Additional Information: full citation, abstract, index terms

We present and evaluate the idea of adaptive processor cache management. Specifically, we describe a novel and general scheme by which we can combine any two cache management algorithms (e.g., LRU, LFU, FIFO, Random) and adaptively switch between them, closely tracking the locality characteristics of a given program. The scheme is inspired by recent work in virtual memory management at the operating system level, which has shown that it is possible to adapt over two replacement policies to provi ...

4 Compilation: Automated compile-time and run-time techniques to increase usable memory in MMU-less embedded systems

Lan S. Bai, Lei Yang, Robert P. Dick

October 2006 Proceedings of the 2006 international conference on Compilers, architecture and synthesis for embedded systems CASES '06

Publisher: ACM Press

Full text available: pdf(1.94 MB) Additional Information: full citation, abstract, references, index terms

Random access memory (RAM) is tightly-constrained in many embedded systems. This is especially true for the least expensive, lowest-power embedded systems, such as sensor network nodes and portable consumer electronics. The most widely-used sensor network nodes have only 4-10 KB of RAM and do not contain memory management units (MMUs). It is very difficult to implement increasingly complex applications under such tight memory constraints. Nonetheless, price and power consumption constraints make ...

**Keywords**: data compression, embedded system, wireless sensor network

5 Full Technical Papers: Learning implicit user interest hierarchy for context in

personalization

Hyoung R. Kim, Philip K. Chan

January 2003 Proceedings of the 8th international conference on Intelligent user interfaces IUI '03

Publisher: ACM Press

Full text available: pdf(191.53 KB)

Additional Information: full citation, abstract, references, citings, index terms

To provide a more robust context for personalization, we desire to extract a continuum of general (long-term) to specific (short-term) interests of a user. Our proposed approach is to learn a user interest hierarchy (UIH) from a set of web pages visited by a user. We devise a divisive hierarchical clustering (DHC) algorithm to group words (topics) into a hierarchy where more general interests are represented by a larger set of words. Each web page can then be assigned to nodes in the hierarchy f ...

Keywords: clustering algorithm, concept clustering, user interest hierarchy, user profile

6 Computation: finite and infinite machines

Marvin L. Minsky January 1967 Book

Publisher: Prentice-Hall, Inc.

Additional Information: full citation, abstract, references, citings, index terms

From the Preface (See Front Matter for full Preface)

Man has within a single generation found himself sharing the world with a strange new species: the computers and computer-like machines. Neither history, nor philosophy, nor common sense will tell us how these machines will affect us, for they do not do "work" as

did machines of the Industrial Revolution. Instead of dealing with materials or energy, we are told that they handle "control" and "information" and even "intellectua ...

### Structured programming

January 1972 Divisible Book Publisher: Academic Press Ltd.

Full text available: pdf(11.44 MB) Additional Information: full citation, abstract, cited by, index terms

In recent years there has been an increasing interest in the art of computer programming, the conceptual tools available for the design of programs, and the prevention of programming oversights and error. The initial outstanding contribution to our understanding of this subject was made by E. W. Dijkstra, whose Notes on Structured Programming form the first and major section of this book. They clearly expound the reflections of a brilliant programmer on the methods which he has hitherto uncon ...

## IBM system/360 principles of operation

**IBM** 

January 1964 Book Publisher: IBM Press

Full text available: pdf(14.82 MB) Additional Information: full citation, abstract, index terms

This manual is a comprehensive presentation of the eharaeteristies, functions, and features of the IBM System/ 360. The material is presented in a direct manner, assuming that the reader has a basic knowledge of IBM data processing systems and has read the IBM System/360 Systems Summary, Form A22-6810. The manual is useful for individual study, as an instruction aid, and as a machine reference manual.

The manual defines System/360 operating principles, central processing unit, instruct ...

# Cryptography and data security

Dorothy Elizabeth Robling Denning

January 1982 Book

Publisher: Addison-Wesley Longman Publishing Co., Inc.

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(19.47 MB)

#### From the Preface (See Front Matter for full Preface)

Electronic computers have evolved from exiguous experimental enterprises in the 1940s to prolific practical data processing systems in the 1980s. As we have come to rely on these systems to process and store data, we have also come to wonder about their ability to protect valuable data.

Data security is the science and study of methods of protecting data in computer and communication systems from unauthorized disclosure ...

## 10 Common knowledge

John Geanakoplos

March 1992 Proceedings of the 4th conference on Theoretical aspects of reasoning about knowledge TARK '92

Publisher: Morgan Kaufmann Publishers Inc.

Full text available: R pdf(3.29 MB) Additional Information: full citation, abstract, references

People, no matter how rational they are, usually act on the basis of incomplete information. If they are rational they recognize their own ignorance and reflect carefully on what they know and what they do not know, before choosing how to act. Furthermore, when rational agents interact, they also think about what the others know, and what the others know about what they know, before choosing how to act. Failing to do so can be disastrous. When the notorious evil genius Professor Moriarty conf ...

11 Memory: Heterogeneous way-size cache

Jaume Abella, Antonio González

June 2006 Proceedings of the 20th annual international conference on Supercomputing ICS '06

Publisher: ACM Press

Full text available: 🔁 pdf(410.54 KB) Additional Information: full citation, abstract, references, index terms

Set-associative cache architectures are commonly used. These caches consist of a number of ways, each of the same size. We have observed that the different ways have very different utilization, which motivates the design of caches with heterogeneous way sizes. This can potentially result in higher performance for the same area, better capabilities to implement dynamically adaptive schemes, and more flexibility for choosing the size of the cache. This paper proposes a novel cache architecture, the ...

Keywords: adaptive, cache memories, low power, set-associative

12 Improving instruction cache performance in OLTP

Stavros Harizopoulos, Anastassia Ailamaki

September 2006 ACM Transactions on Database Systems (TODS), Volume 31 Issue 3

Publisher: ACM Press

Full text available: 🔁 pdf(783.16 KB) Additional Information: full citation, abstract, references, index terms

Instruction-cache misses account for up to 40% of execution time in online transaction processing (OLTP) database workloads. In contrast to data cache misses, instruction misses cannot be overlapped with out-of-order execution. Chip design limitations do not allow increases in the size or associativity of instruction caches that would help reduce misses. On the contrary, the effective instruction cache size is expected to further decrease with the adoption of multicore and multithreading ...

Keywords: Instruction cache, cache misses

13 Writing efficient programs

Jon Louis Bentley January 1982 Book

Publisher: Prentice-Hall, Inc.

Additional Information: full citation, abstract, references, citings, index terms

The primary task of software engineers is the cost-effective development of maintainable and useful software. There are many secondary problems lurking in that definition. One such problem arises from the term "useful": to be useful in the application at hand, software must often be efficient (that is, use little time or space). The problem we will consider in this book is building efficient software systems.

There are a number of levels at which we may confront the problem of efficien ...

14 Chapter II: Notes on data structuring

C. A. R. Hoare

January 1972 Structured programming

Publisher: Academic Press Ltd.

Full text available: pdf(4.64 MB) Additional Information: full citation, references

## Optimal clip ordering for multi-clip queries

Raymond T. Ng, Paul Shum

December 1998 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 7 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: 📆 pdf(122.95 KB) Additional Information: full citation, abstract, index terms

A multi-clip query requests multiple video clips be returned as the answer of the query. In many applications and situations, the order in which these clips are to be delivered does not matter that much to the user. This allows the system ample opportunities to optimize system throughput by using schedules that maximize the effect of piggybacking. In this paper, we study how to find such optimal schedules. In particular, we consider two optimization criteria: (i) one based on maximizing the numb ...

Keywords: Admission control, Bipartite graph matching, Performance of multimedia systems

### 16 DataScalar architectures

Doug Burger, Stefanos Kaxiras, James R. Goodman

May 1997 ACM SIGARCH Computer Architecture News, Proceedings of the 24th annual international symposium on Computer architecture ISCA '97, Volume 25 Issue 2

**Publisher: ACM Press** 

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(2.11 MB) terms

DataScalar architectures improve memory system performance by running computation redundantly across multiple processors, which are each tightly coupled with an associated memory. The program data set (and/or text) is distributed across these memories. In this execution model, each processor broadcasts operands it loads from its local memory to all other units. In this paper, we describe the benefits, costs, and problems associated with the DataScalar model. We also present simulation results of ...

# 17 Architectural support for scalable speculative parallelization in shared-memory

<u>multiprocessors</u>

Marcelo Cintra, José F. Martínez, Josep Torrellas

May 2000 ACM SIGARCH Computer Architecture News, Proceedings of the 27th annual international symposium on Computer architecture ISCA '00, Volume 28 Issue 2

Publisher: ACM Press

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> Full text available: pdf(253.29 KB) <u>terms</u>

Speculative parallelization aggressively executes in parallel codes that cannot be fully parallelized by the compiler. Past proposals of hardware schemes have mostly focused on single-chip multiprocessors (CMPs), whose effectiveness is necessarily limited by their small size. Very few schemes have attempted this technique in the context of scalable shared-memory systems. In this paper, we present and evaluate a new hardware scheme for scalable speculative parallelization. This de ...

# 18 A unified architectural tradeoff methodology

C.-H. Chen, A. K. Somani

April 1994 ACM SIGARCH Computer Architecture News, Proceedings of the 21ST annual international symposium on Computer architecture ISCA '94, Volume 22 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available: pdf(1.06 MB)

Additional Information: full citation, abstract, references, citings, index

We present a unified approach to assess the trade-off of architecture techniques that affect mean memory access time. The architectural features we consider include cache hit ratio, processor stalling features, line size, memory cycle time, the external data bus width of a processor, pipelined memory system, and read by-passing write buffers. We demonstrate how each of these features can be traded off to achieve the desired performance. The performance of an architecture feature is quantified in ...

19 TinyDB: an acquisitional query processing system for sensor networks



Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong March 2005 ACM Transactions on Database Systems (TODS), Volume 30 Issue 1

Publisher: ACM Press

Full text available: pdf(1.67 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

We discuss the design of an acquisitional query processor for data collection in sensor networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (sampled) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the a priori existence of data. We discuss simple extensions to SQL for controlli ...

**Keywords**: Query processing, data acquisition, sensor networks

20 Adaptive, fine-grained sharing in a client-server OODBMS: a callback-based



approach

Markos Zaharioudakis, Michael J. Carey, Michael J. Franklin December 1997 ACM Transactions on Database Systems (TODS), Volume 22 Issue 4

Publisher: ACM Press

Full text available: pdf(441.80 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

For reasons of simplicity and communication efficiency, a number of existing objectoriented database management systems are based on page server architectures; data pages are their minimum unit of transfer and client caching. Despite their efficiency, page servers are often criticized as being too retrictive when it comes to concurrency, as existing systems use pages as the minimum locking unit as well. In this paper we show how to support object-level locking in a page-server context. Sev ...

Keywords: cache coherency, cache consistency, client-server databased, fine-grained sharing, object-oriented databases, performance analysis

Results 1 - 20 of 200

Result page: 1 2 3 4 5 6 7 8 9 10

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player

# **EAST Search History**

| Ref . | Hits | Search Query  | DBs   | Default<br>Operator | Plurals | Time Stamp       |
|-------|------|---|---|---------------------|---------|------------------|
| L1    | 1    | (subordinat\$4 or sub-ordinat\$4)<br>same index\$4 same temporar\$4<br>same map\$4  | US-PGPUB;<br>USPAT;<br>EPO; JPO;<br>DERWENT;<br>IBM_TDB | OR                  | ON ·    | 2007/06/06 16:35 |
| L2    | 1    | predetermin\$4 same threshol\$4<br>same exceed\$4 same check\$4<br>same (subordinat\$4 or<br>sub-ordinat\$4) same cach\$4 | US-PGPUB;<br>USPAT;<br>EPO; JPO;<br>DERWENT;<br>IBM_TDB | OR                  | ON      | 2007/06/06 16:37 |
| L3    | 6    | cach\$4 same elaps\$4 same tim\$4 same check\$4 same count\$4 same period\$1  | US-PGPUB;<br>USPAT;<br>EPO; JPO;<br>DERWENT;<br>IBM_TDB | OR                  | ON      | 2007/06/06 16:38 |
| L4    | 1    | cach\$4 same elaps\$4 same tim\$4 same check\$4 same count\$4 same period\$1 same predetermi\$4                           | US-PGPUB;<br>USPAT;<br>EPO; JPO;<br>DERWENT;<br>IBM_TDB | OR                  | ON      | 2007/06/06 16:39 |